

Local Optimality of User Choices and Collaborative Competitive Filtering

Shuang Hong Yang
 College of Computing
 Georgia Institute of Technology
 Atlanta, GA 30332
 shy@gatech.edu

Abstract

While a user's preference is directly reflected in the interactive choice process between her and the recommender, this wealth of information was not fully exploited for learning recommender models. In particular, existing collaborative filtering (CF) approaches take into account only the binary events of user actions but totally disregard the contexts in which users' decisions are made. In this paper, we propose *Collaborative Competitive Filtering* (CCF), a framework for learning user preferences by modeling the choice process in recommender systems. CCF employs a multiplicative latent factor model to characterize the dyadic utility function. But unlike CF, CCF models the user behavior of choices by encoding a local competition effect. In this way, CCF allows us to leverage dyadic data that was previously lumped together with missing data in existing CF models. We present two formulations and an efficient large scale optimization algorithm. Experiments on three real-world recommendation data sets demonstrate that CCF significantly outperforms standard CF approaches in both offline and online evaluations.

1 Introduction

Recommender systems have become a core component for today's personalized online businesses [23, 9]. With the abilities of connecting various items (e.g., retailing products, movies, News articles, advertisements, experts) to potentially interested users, recommender systems enable online webshops (e.g. Amazon, Netflix, Yahoo!) to expand the marketing efforts from historically a few best-sellings toward a large variety of long-tail (niche) products [4, 9, 28]. Such abilities are

endowed by a personalization algorithm for identifying the preference of each individual user, which is at the heart of a recommender system.

Predicting user preference is challenging. Usually, the user and item spaces are very large yet the observations are extremely sparse. Learning from such rare, noisy and largely missing evidences has a high risk of overfitting. Indeed, this *data sparseness* issue has been widely recognized as a critical challenge for constructing effective recommender systems.

A straightforward way for building recommender would be to learn a user’s preference based on the prior interactions between her and the recommender system. Typically, such interaction is an “opportunity give-and-take” process (c.f. Table 1), where at each interaction:

- 1) a user u inquires the system (e.g. visits a movie recommendation web site);
- 2) the system offers a set of (personalized) opportunities (i.e. items) $\mathcal{O} = \{i_1, \dots, i_l\}$ (e.g. recommends a list of movies of potential interest to the user);
- 3) the user chooses one item $i^* \in \mathcal{O}$ (or more) from these offers and takes actions accordingly (e.g. click a link, rent a movie, view a News article, purchase a product).

Somewhat surprisingly, this interaction process has not been fully-exploited for learning recommenders. Instead, research on recommender systems has focused almost exclusively on recovering user preference by completing the matrix of user actions (u, i^*) while the actual contexts in which user decisions are made are totally disregarded. In particular, Collaborative Filtering (CF) approaches only captures the action dyads (u, i^*) while the contextual dyads (i.e. $\{(u, i)\}$ for all $i \in \mathcal{O}$ and $i \neq i^*$) are typically treated as missing data. For example, the rating-oriented models aim to approximating the ratings that users assigned to items [25, 20, 24, 1, 6, 15]; the recently proposed ranking-oriented algorithms [29, 16] attempt to recover the ordinal ranking information derived from the ratings. Although this formulation of the recommendation problem has led to numerous algorithms which excel at a number of data sets, including the prize-winning work of [15], we argue here that the formulation is inherently flawed — a preference for *Die Hard* given a generic set of movies only tells us that the user appreciates action movies; however, a preference for *Die Hard* over *Terminator* or *Rocky* suggests that the user might favor Bruce Willis over other action heroes. In other words, the context of user choice is vital when estimating user preferences.

When it comes to modeling of user-recommender interactions, an important question arises: what is the fundamental mechanism underlying the user choice behaviors? As reflected by its name, collaborative filtering is based on the notion of “collaboration effects” that *similar items get similar responses from similar users*. This assumption is essential because by encoding the “collaboration” among users

User	Offer set	Choice
u_1	$[i_1, i_2, i_3, i_5]$	i_2
u_2	$[i_2, i_3, i_4, i_5]$	i_2
u_3	$[i_1, i_3, i_5, i_6]$	i_5
u_4	$[i_2, i_3, i_4, i_6]$	i_3
u_5	$[i_1, i_3, i_4, i_5]$	i_4
u_6	$[i_1, i_4, i_5, i_6]$	i_6

	i_1	i_2	i_3	i_4	i_5	i_6
u_1	.	1
u_2	.	1
u_3	1	.
u_4	.	.	1	.	.	.
u_5	.	.	.	1	.	.
u_6	1

Table 1: An example of user-recommender interactions and the derived observation matrix: entries with value “1” denote the action dyads; dyads that are observed without user actions (e.g. offered by the recommender but not picked by the user) are marked with dots (“.”). CF trains only on the 1 entries while the . entries are treated as missing data. CCF distinguishes between unseen entries and entries marked with dots.

or among items or both, CF greatly alleviates the issue of data sparseness and in turn makes more reliable predictions based on the somewhat *pooled* evidences across different items/users.

It has long been recognized in psychology and economics that, besides the effect of collaboration [5, 21], another mechanism governs users’ behavior — *competition* [17, 19, 3]. In particular, items turn to compete with each other for the attention of users; therefore, axiomatically, user u will pick the best item i^* (i.e. the one with highest utility) when confronted by the set of alternatives \mathcal{O} . For example, consider a user with a penchant for action movies by Arnold Schwarzenegger. Given the choice between *Sleepless in Seattle* and *Die Hard* he will likely choose the latter. However, when afforded the choice between the oeuvres of Schwarzenegger, Diesel or Willis, he’s clearly more likely to choose Schwarzenegger over the works of Willis. To capture user’s preference more accurately, it is therefore *essential* for a recommender model to take into account such local competition effect. Unfortunately, this effect is absent in a large number of collaborative filtering approaches.

In this paper, we present Competitive Collaborative Filtering (CCF) for learning recommender models by modeling users’ choice behavior in their interactions with the recommender system. Similar to matrix factorization approaches for CF, we employ a multiplicative latent factor model to characterize the dyadic utility function (i.e. the utility of an item to a user). In this way, CCF encodes the collaboration effect among users and items similar to CF. But instead of learning only the action dyads (i.e. (u, i^*) or the “1” entries in Table 1), CCF bases the factorization learning on the whole user-recommender interaction traces. It therefore leverages not only the action dyads (u, i^*) but also the dyads in the context without user actions (i.e. (u, i) for all $i \in \mathcal{O}$ and $i \neq i^*$ or the dot entries in Table 1), which were treated as potentially missing data in CF approaches.

To leverage the entire interaction trace for latent factor learning, we devise

probabilistic models or optimization objectives to encode the local competition effect underlying the user choice process. We present two formulations with different flavors. The first formulation is derived from the *multinomial logit model* that has been widely used for modeling user choice behavior (e.g. choice of brands) in psychology [17], economics [18, 19] and marketing science [11]. The second formulation relates closely to the ordinal regression models in content filtering [12] (e.g. web search ranking). Essentially, both formulations attempt to encode “local optimality of user choices” to encourage that every opportunity i^* taken by a user u be locally the best in the context of the opportunities \mathcal{O} offered to her. From a machine learning viewpoint, CCF is a hybrid of *local* and *global* learning, where a global matrix factorization model is learned by optimizing a local context-aware loss function. We discuss the implementation of CCF, establish efficient learning algorithms and deliver an package that allows distributed optimization on streaming data.

Experiments were conducted on three real-world recommendation data sets. First, on two dyadic data sets, we show that CCF improves over standard CF models by up to 50+% in terms of offline top- k ranking. Furthermore, on a commercial recommender system, we show that CCF significantly outperform CF models in both offline and online evaluations. In particular, CCF achieves up to 7% improvement in offline top- k ranking and up to 13% in terms of online click rate prediction.

Outline: §2 describes the problem formulation, the backgrounds and motivates CCF. §3 presents the detailed CCF models, learning algorithms and our distributed implementation. §4 reports experiments and results. §5 reviews related work and §6 summarizes the results.

2 Preliminaries

2.1 Problem formulation

Consider the user-system interaction in a recommender system: we have users $u \in \mathcal{U} := \{1, 2, \dots, U\}$ and items $i \in \mathcal{I} = \{1, 2, \dots, I\}$; when a user u visits the site, the system recommends a set of items $\mathcal{O} = \{i_1, \dots, i_l\}$ and u in turn chooses a (possibly empty) subset $\mathcal{D} \subseteq \mathcal{O}$ from \mathcal{O} and takes actions accordingly (e.g. buys some of the recommended products). For ease of explanation, let us temporarily assume $\mathcal{D} = \{i^*\}$, i.e. \mathcal{D} is not empty and contains exactly one item i^* . More general scenarios shall be discussed later.

To build the recommender system, we record a collection of historical interactions in the form of $\{(u_t, \mathcal{O}_t, \mathcal{D}_t)\}$, where t is the index of a particular interaction session. Our goal is to generate recommendations $\mathcal{O}_{\tilde{t}}$ for an incoming visit \tilde{t} of

user u_i such that the user’s satisfaction is maximized. Hereafter, we refer to \mathcal{U} as *user space*, \mathcal{I} as *item space*, \mathcal{O}_t as *offer set* or *context*, \mathcal{D}_t as *decision set*, and i^* as a *decision*.

A key component of a recommender system is a model $r(u, i)$ that characterizes the utility of an item $i \in \mathcal{I}$ to a user $u \in \mathcal{U}$, upon which recommendations for a new inquiry from user u could be done by simply ranking items based on $r(u, i)$ and recommending the top-ranked ones. Collaborative filtering is by far the most well-known method for modeling such dyadic responses.

2.2 Collaborative filtering

In collaborative filtering we are given observations of dyadic responses $\{(u, i, y_{ui})\}$ with each y_{ui} being an observed response (e.g. user’s rating to an item, or indication of whether user u took an action on item i). The whole mapping:

$$(u, i) \rightarrow y_{ui} \text{ where } u \in \mathcal{U}, i \in \mathcal{I}$$

constitutes a large matrix $Y \in \mathcal{Y}^{|\mathcal{U}| \times |\mathcal{I}|}$. While we might have millions of users and items, only a tiny proportion (considerably less than 1% in realistic datasets) of entries are observable. Note the subtle difference in terms of the data representation: while we record entire sessions, CF only records the dyadic responses.

Collaborative filtering explores the notion of “collaboration effects”, i.e., similar users have similar preference to similar items. By encoding collaboration, CF pools the sparse observations in such a way that for predicting $r(u, i)$ it also borrows observations from other users/items. Generally speaking, existing CF methods fall into either of the following two categories.

Neighborhood models. A popular class of approaches to CF is based on propagating the observations of responses among items or users that are considered as neighbors. The model first defines a similarity measure between items / users. Then, an unseen response between user u and item i is approximated based on the responses of neighboring users or items [25, 20], for example, by simply averaging the neighboring responses with similarities as weights.

Latent factor models. This class of methods learn predictive latent factors to estimate the missing dyadic responses. The basic idea is to associate latent factors¹, $\phi_u \in \mathbb{R}^k$ for each user u and $\phi_i \in \mathbb{R}^k$ for each item i , and assume a multiplicative model for the dyadic response,

$$p(y_{ui}|u, i) = p(y_{ui}|r_{ui}; \Theta),$$

¹Throughout this paper, we assume each latent factor ϕ contains a constant component so as to absorb user/item-specific offset into latent factors.

where Θ denotes the set of hyper-parameters, the utility is assumed as a multiplicative function of the latent factors,

$$r(u, i) = \phi_u^\top \phi_i.$$

This way the factors could explain past responses and in turn make prediction for future ones. This model implicitly encodes the Aldous-Hoover theorem [13] for exchangeable matrices – y_{ui} are independent of each other given ϕ_u and ϕ_i . In essence, it amounts to a low-rank approximation of the matrix Y that naturally embeds both users and items into a vector space in which the distances directly reflect the semantic relatedness.

To design a concrete model [2, 22, 26], one needs to specify a distribution for the dependence. Afterwards, the model boils down to an optimization problem. For example two commonly-used formulations are:

- **ℓ_2 regression** The most popular learning formulation is to minimize the ℓ_2 loss within an empirical risk minimization framework [14, 24]:

$$\min_{\phi} \sum_{(u,i) \in \Omega} (y_{ui} - \phi_u^\top \phi_i)^2 + \lambda_{\mathcal{U}} \sum_{u \in \mathcal{U}} \|\phi_u\|^2 + \lambda_{\mathcal{I}} \sum_{i \in \mathcal{I}} \|\phi_i\|^2,$$

where Ω denotes the set of (u, i) dyads for which the responses y_{ui} are observed, $\lambda_{\mathcal{U}}$ and $\lambda_{\mathcal{I}}$ are regularization weights.

- **Logistic** Another popular formulation [22, 1] is to use logistic regression by optimizing the cross-entropy

$$\begin{aligned} \min_{\phi} \sum_{(u,i) \in \Omega} \log \left[1 + \exp(-\phi_u^\top \phi_i) \right] \\ + \lambda_{\mathcal{U}} \sum_{u \in \mathcal{U}} \|\phi_u\|^2 + \lambda_{\mathcal{I}} \sum_{i \in \mathcal{I}} \|\phi_i\|^2 \end{aligned}$$

2.3 Motivating discussions

Collaborative filtering approaches have made substantial progresses and are currently the state-of-the-art techniques for recommender system. However, we argue here that CF approaches might be a bit lacking in several aspects. First of all, although data sparseness is a big issue, CF does not fully leverage the wealth of user behavior data. Take the user-recommender interaction process described in §2.1 as an example (c.f. Table 1), CF methods typically use only the action dyad (u, i^*) of each session while other dyads $\{(u, i) | i \in \mathcal{O}, i \neq i^*\}$ are treated missing and

totally disregarded, which could be wasteful of the invaluable learning resource because these non-action dyads are not totally useless, as shown by the experiments in this paper.

Secondly, most existing CF approaches learn user preference collaboratively by either approximating the dyadic responses $\{y_{ui}^*\}$ [25, 20, 24, 1, 6, 15] or preserving the ordinal ranking information derived from the dyadic responses [29, 16]; none of them models the user choice behavior in recommender systems. Particularly, as users choose from competing alternatives, there is naturally a local competition effect among items being offered in a session. Our work show that this effect could be an important clue for learning user preference.

Because latent factor models are very flexible and could be under-determined (or over-parameterized) even for rather moderate number of users/items. With the above two limitations, CF approaches are vulnerable to over-fitting [1, 15]. Particularly, while most existing CF models might learn consistently on user ratings (numerical value typically with five levels) if given enough training data, they usually perform poorly on binary responses. For example, for the aforementioned interaction process (*c.f.* Table 1), the response y_{ui} is typically a binary event indicating whether or not item i was accepted by the user u . With the non-action dyads being ignored, the responses are exclusively positive observations (either $y_{ui} = 1$ or missing). As a result, we will obtain an overly-optimistic estimator that biases toward positive responses and predicts positive for almost all the incoming dyads (See §4.1 for empirical evidences).

3 Collaborative competitive filtering

We present a novel framework for recommender learning by modeling the system-user interaction process. The key insight is that the contexts \mathcal{O}_t in which user's decisions are made should be taken into account when learning recommender models. In practice, a user u could make different decisions when facing different contexts \mathcal{O}_t . For instance, an item i would not have been chosen by u if it were not presented to her at the first place; likewise, user u could choose another item if the context \mathcal{O}_t changes such that a better offer (e.g., a more interesting item) is presented to her.

In this section, we describe the framework of collaborative-competitive filtering. We start with some axiomatic views of the user choice behaviors. Following that, we present the learning formulation of CCF. We then develop the optimization algorithms and implementation techniques. We close the section with a discussion of useful extensions.

3.1 Local optimality of user choices

Formally, the individual choice process (i.e. user-recommender interactions) in a recommender system can be viewed as an instance of the *opportunity give-and-take* (GAT) process.

DEFINITION [GAT]: *An opportunity give-and-take process is a process of interactions among an agent u , a system S and a set of opportunities \mathcal{I} ; at an interaction t :*

- u is given a set of opportunities $\mathcal{O}_t \subset \mathcal{I}$ by S ;
- u makes the decision by taking one of the opportunities: $i_t^* \in \mathcal{O}_t$;
- Each opportunity $i \in \mathcal{O}_t$ could potentially give u a revenue (utility) of r_{ui} if being taken or 0 otherwise.

Note that we assume the agent is *a priori* not aware of all the items, and only through the recommender S can she get to know the items, therefore other items that are not in \mathcal{O}_t is inaccessible to u at interaction t . This is reasonable considering that the number of item is usually very large. Moreover, we assume an agent u is a rational decision maker: she knows that her choice of item i will be at the expense of others $i' \in \mathcal{O}_t$, therefore she compares among alternatives before making her choice. In other words, for each decision, u considers both **revenue** and **opportunity cost**, and decides which opportunity to take based on the potential **profit** of each opportunity in \mathcal{O} . Specifically, the opportunity cost c_{ui} is the potential loss of u from taking an opportunity i that excludes her to take other opportunities: $c_{ui} = \max\{r_{ui'} : i' \in \mathcal{O} \setminus i\}$; the profit $\pi_{ui} = r_{ui} - c_{ui}$ is the net gain of an decision. By drawing the rational decision theory [17], we present the following principle of individual choice behavior.

PROPOSITION: *A rational decision is a decision maximizing the profit: $i^* = \arg \max_{i \in \mathcal{O}} \pi_{ui}$.*

This proposition implies the constraint of “local optimality of user choice”, a local competitive effect restricting that the agent u always chooses the offer that is locally optimal in the context of the offer set \mathcal{O}_t .

3.2 Collaborative competitive filtering

The local-optimality principle induces a constraint which could be translated to an objective function for recommender learning:

$$\begin{aligned} & \forall i^* \in \mathcal{D}_t, \quad r_{ui^*} \geq \max\{r_{ui} | i \in \mathcal{O}_t \setminus \mathcal{D}_t\} \\ \text{or } & P(i^* \text{ is taken}) = P(r_{ui^*} \geq \max\{r_{ui} | i \in \mathcal{O}_t \setminus \mathcal{D}_t\}). \end{aligned} \quad (1)$$

This objective is, however, problematic. First, the inequality constraint restricts the utility function only up to an arbitrary order-preserving transformation (e.g. a monotonically increasing function), and hence cannot yield a unique solution (e.g. point estimation) [18]. Second, optimization based on the induced objective is computationally intractable due to the \max operator. To this end, we present two surrogate objectives, which both are computationally efficient and show close connections to existing models.

3.2.1 Softmax model

Our first formulation is based on the random utility theory [17, 18] which has been extensively used for modeling choice behavior in economics [19] and marketing science [11]. In particular, we assume the utility function consists of two components $r_{ui} + e_{ui}$, where: (1) r_{ui} is a deterministic function characterizing the intrinsic interest of user u to item i , for which we use the latent factor model to quantify $r_{ui} = \phi_u^\top \phi_i$; (2) the second part e_{ui} is a stochastic error term reflecting the uncertainty and complexness of the choice process². Furthermore, we assume the error term e_{ui} is an independently and identically distributed Weibull (extreme point) variable:

$$\Pr(e_{ui} \leq \epsilon) = e^{-e^{-\epsilon}}.$$

Together with the local-optimality principle, these two constraints yield the following *multinomial logit model* [19, 18, 11]:

$$p(i^* = i | u, \mathcal{O}) = \frac{e^{r_{ui}}}{\sum_{j \in \mathcal{O}} e^{r_{uj}}} \text{ for all } i \in \mathcal{O}. \quad (2)$$

Intuitively, this model enforces the local-optimality constraint by using the *softmax* function as a surrogate of *max*.

Given a collection of training interactions $\{(u_t, \mathcal{O}_t, i_t^*)\}$, the latent factors can be estimated using penalized maximum likelihood via

$$\begin{aligned} \min_{\phi} \quad & \sum_t \log \left[\sum_{i \in \mathcal{O}_t} \exp(\phi_{u_t}^\top \phi_i) \right] - \phi_{u_t}^\top \phi_{i_t^*} \\ & + \lambda_{\mathcal{U}} \sum_{u \in \mathcal{U}} \|\phi_u\|^2 + \lambda_{\mathcal{I}} \sum_{i \in \mathcal{I}} \|\phi_i\|^2. \end{aligned} \quad (3)$$

While the above formulation is a convex optimization w.r.t. r_{ui} as each of the objective terms in Eq.(3) is strongly concave, it is nonconvex w.r.t. the latent factors ϕ . We postpone the discussion of optimization algorithms to §3.3.

²The error term essentially accounts for all the subtle, uncertain and unmeasurable factors that influence user choice behaviors, for example, a user's mood, past experience, or other factors (e.g., whether the decision is made in a hurry, together with her friends, or totally unconsciously)

3.2.2 Hinge model

Our second formulation is based on a simple reduction of the local-optimality constraint. Note that, from Eq(1), it follows that:

$$\begin{aligned} P(i = i^* | u, \mathcal{O}) &= P((r_{ui^*} - r_{ui}) > (e_{ui} - e_{ui^*}), \forall i \in \mathcal{O}) \\ &\leq P((r_{ui^*} - \bar{r}_{ui}) > (\bar{e}_{ui} - e_{ui^*})), \end{aligned}$$

where $\bar{r}_{ui} = \frac{1}{|\mathcal{O}|-1} \sum_{i \in \mathcal{O} \setminus i^*} r_{ui}$ is the average potential utility that u could possibly gain from the non-chosen items. Intuitively, the above model encourages that the utility difference between choice and non-chosen items, $r_{ui^*} - \bar{r}_{ui}$, to be non-trivially greater than random errors. Based on this notion, we present the following formulation which views the task as a pairwise preference learning problem [12] and uses the non-choices averagely as negative preferences.

$$\begin{aligned} \min_{\theta, \xi} \quad & \sum_t \xi_t + \lambda_{\mathcal{U}} \sum_{u \in \mathcal{U}} \|\phi_u\|^2 + \lambda_{\mathcal{I}} \sum_{i \in \mathcal{I}} \|\phi_i\|^2 \\ \text{s.t.} \quad & r_{ui_t^*} - \frac{1}{|\mathcal{O}_t|-1} \sum_{i \in \mathcal{O}_t \setminus \{i_t^*\}} r_{ui} \geq 1 - \xi_t \text{ and } \xi_t \geq 0. \end{aligned} \quad (4)$$

This formulation is directly related to the maximum score estimation [18] of the multinomial logit model Eq(2). Intuitively, it directly reflects the insight that user decisions are usually made by comparing alternatives and considering the *difference* of potential utilities. In other words, it learns latent factors by maximizing the marginal utility between user choice and the average of non-choices.

Again, the optimization is convex w.r.t. r_{ui} , but nonconvex w.r.t. the latent factors, therefore the standard optimization tools such as the large variety of RankSVM [12] solvers are not directly applicable.

3.2.3 Complexity

It is worth noting that our CCF formulations have an appealing linear complexity, $O(|\mathcal{I}| \times |\mathcal{O}|)$, where the offer size $|\mathcal{O}|$ is typically a very small number. For example, Netflix recommends $|\mathcal{O}| = 7$ movies for each visit, and Yahoo! frontpage highlights $|\mathcal{O}| = 4$ hot news for each browser. Therefore, CCF has the same-order complexity as the rating-oriented CF models. Note that the ranking-oriented CF approaches [29, 16] are much more expensive – for each user u , the learning complexity is quadratic $O(|\mathcal{I}|^2)$ as they learn preference of each user by comparing every pair of the items.

3.3 Learning algorithms

As we have already mentioned, due to the use of bilinear terms, both of the two CCF variants are nonconvex optimization problems regardless of the choice of the loss functions. While there *are* convex reformulations for some settings they tend to be computationally inefficient for large scale problems as they occur in industry — the convex formulations require the manipulation of a full matrix which is impractical for anything beyond thousands of users.

Moreover, the interactions between user and items change over time and it is desirable to have algorithms which process this information incrementally. This calls for learning algorithms that are sufficiently efficient and preferably capable to update dynamically so as to reflect upcoming data streams, therefore excluding offline learning algorithms such as classical SVD-based factorization algorithms [15] or spectral eigenvalue decomposition methods [16] that involve large-scale matrices.

We use a distributed stochastic gradient variant with averaging based on the Hadoop MapReduce framework. The basic idea is to decompose the objectives in (3) or (4) by running stochastic optimization on sub-blocks of the interaction traces in parallel in the Map phase, and to combine the results for ϕ_i in the Reduce phase. The basic structure is analogous to [6, 32].

Stochastic Optimization. We derive a stochastic gradient descent algorithm to solve the optimization described in Eq(3) or Eq(4). The algorithm is computationally efficient and decouplable among different interactions and users, therefore amenable for parallel implementation.

The algorithm loops over all the observations and updates the parameters by moving in the direction defined by negative gradient. Specifically, we can carry out the following update equations on each machine separately:

- For all $i \in \mathcal{O}_t$ do $\phi_i \leftarrow \phi_i - \eta [l'(\phi_u^\top \phi_i) \phi_u + \lambda_{\mathcal{I}} \phi_i]$.
- For each u do $\phi_u \leftarrow \phi_u - \eta [\sum_{i \in \mathcal{O}_t} l'(\phi_u^\top \phi_i) \phi_i + \lambda_{\mathcal{U}} \phi_u]$.

Here η is the learning rate³. The gradients are given by:

$$l'_{\text{Softmax}}(r_{ui}) = \frac{\exp(r_{ui})}{\sum_{j \in \mathcal{O}} \exp(r_{uj})} - \delta_{i,i^*} \quad (5)$$

$$l'_{\text{Hinge}}(r_{ui}) = -\frac{|\mathcal{O}| \delta_{i,i^*} - 1}{|\mathcal{O}| - 1} H(1 - r_{ui^*} + \bar{r}_{ui}) \quad (6)$$

³We carry out an annealing procedure to discount η by a factor of 0.9 after each iteration, as suggested by [14].

where $H(\cdot)$ is the Heaviside function, i.e. $H(x) = 1$ if $x > 0$ and $H(x) = 0$ otherwise.⁴

Feature Hashing. A key challenge in learning CCF models on large-scale data is that the storage of parameters as well as observable features requires a large amount of memory and a reverse index to map user IDs to memory locations. In particular in recommender systems with hundreds of millions of users the memory requirement would easily exceed what is available on today’s computers (100 million users with 100 latent feature dimensions each amounts to 40GB of RAM). We address this problem by implementing feature hashing [30] on the space of matrix elements. In particular, by allowing random collisions and applying hash mapping to the latent factors (i.e. ϕ), we keep the entire representation in memory, thus greatly accelerating optimization.

3.4 Extensions

We now discuss two extensions of CCF to address the fact that in some cases users choose not to respond to an offer at all and that moreover we may have observed features in addition to the latent representation discussed so far.

Sessions without response

In establishing the CCF framework for modeling the user choice behavioral data, we assumed that for each user-system interaction t , the decision set \mathcal{D}_t contains at least one item. This assumption is, however, not true in practice. A user’s visit at a recommender system does not always yields an action. For example, users frequently visit online e-commerce website without making any purchase, or browse a news portal without clicking on an ad. Actually, such nonresponded visits may account for a vast majority of the traffics that an recommender system receives. Moreover, different users may have different propensities for taking an action. Here, we extend the multinomial logit model to modeling both responded and nonresponded interactions, $(u_t, \mathcal{O}_t, i_t^*)$ and $(u_t, \mathcal{O}_t, \emptyset)$ respectively.

This is accomplished by adding a scalar θ_u for each user u to capture the *action threshold* of user u . We assume that, at an interaction t , user u_t takes an effective action only if she feels that the overall quality of the offers \mathcal{O}_t are good enough and worth the spending of her attention. In keeping with the logistic model this means

⁴In our implementation, we approximate this by the continuous function $\frac{1}{1+e^{-100x}}$. This helps with convergence.

that

$$p(i^* = i|u, \mathcal{O}) = \frac{\exp(\phi_u^\top \phi_i)}{\exp(\theta_u) + \sum_{j \in \mathcal{O}} \exp(\phi_u^\top \phi_j)} \quad (7)$$

for all $i \in \mathcal{O}$ and the probability of no response is given by the remainder, that is by $\frac{\exp(\theta_u)}{\exp(\theta_u) + \sum_{j \in \mathcal{O}} \exp(\phi_u^\top \phi_j)}$. In essence, this amounts to a model where the ‘non-response’ has a certain reserve utility that needs to be exceeded for a user to respond. We may extend the hinge model in the same spirit (we use a trade-off constant $C > 0$ to calibrate the importance of the non-responses).

$$\begin{aligned} \min_{\phi, \xi, \epsilon, \theta} \quad & \sum_t \xi_t + C \sum_t \epsilon_t + \lambda_{\mathcal{U}} \sum_{u \in \mathcal{U}} \|\phi_u\|^2 + \lambda_{\mathcal{I}} \sum_{i \in \mathcal{I}} \|\phi_i\|^2 \\ \text{subject to} \quad & r_{u_t i_t^*} - \bar{r}_{u_t \tilde{i}} - \theta_u \geq 1 - \xi_t \text{ for all } i^* \in \mathcal{D}_t \text{ if } \mathcal{D}_t \neq \emptyset \\ & \theta_u - r_{u_t i} \geq 1 - \epsilon_t, \quad \forall i \in \mathcal{O}_t \text{ if } \mathcal{D}_t = \emptyset \\ & r_{ui} = \phi_u^\top \phi_i \text{ and } \xi_t \geq 0 \text{ and } \epsilon_t \geq 0 \end{aligned} \quad (8)$$

Table 2: Statistics of data sets.

	#user	#item	#dyads	offer size
Social	1.2M	400	29M	-
Netflix-5star	0.48M	18K	100M	-
News	3.6M	2.5K	110M	4

Content features

In previous sections, we use a plain latent factor model for quantifying utility, i.e. $r_{ui} = \phi_u^\top \phi_i$. A known drawback [1] of such model is that it only captures dyadic data (responses), and therefore generalizes poorly for completely new entities, i.e. unseen users or items, of which the observations are missing at the training stage. Here, we extend the model by incorporating content features. In particular, we assume that, in addition to the latent features ϕ s, there exist some observable properties $x_u \in \mathbb{R}^m$ (e.g. a user’s self-crafted registration files) for each user u , and $x_i \in \mathbb{R}^n$ (e.g. a textual description of an item) for each item i . We then assume the utility r_{ui} as a function of both types of features (i.e. observable and latent):

$$r_{ui} \sim p(r_{ui} | \phi_u^\top \phi_i + x_u^\top M x_i; \Theta)$$

where the matrix $M \in \mathbb{R}^{m \times n}$ provides a bilinear form for characterizing the utility based on the content features of the corresponding dyads. This model integrates both collaborative filtering [15] and content filtering [7]. On the one hand, if the user u or item i has no or merely non-informative observable features, the model degrades to a factorization-style utility model [24]. On the other hand, if we assume that ϕ_u and ϕ_i are irrelevant, for instance, if i or j is totally new to the system such that there is no interaction involving either of them as in a cold-start setting, this model becomes the classical content-based relevance model commonly used in, e.g. webpage ranking [31], advertisement targeting [6], and content recommendation [7].

4 Experiments

We report experimental results on two test-beds. First, we evaluate the CCF models with CF baselines on two dyadic data sets with simulated choice contexts. The choice of simulated data generated from CF datasets was made since we are unaware of any *publicly* available datasets directly suitable for CCF. Furthermore, we extend our evaluation to a more strict setting based on user-system interaction session data from a commercial recommender system.

4.1 Dyadic response data

We use dyadic data with binary responses, i.e. $\{(u, i, y_{ui})\}$ where $y_{ui} \in \{1, \text{missing}\}$. We compare different recommender models in terms of their top- k ranking performance.

Social network data. The first data set we used was collected from a commercial social network site, where a user expresses her preference for an item with an explicit indication of “like”. We examine data collected for about one year, involving hundreds of millions of users and a large collection of applications, such as games, sports, news feeds, finance, entertainment, travel, shopping, and local information services. Our evaluation focuses on a random subset consisting of about 400 items, 1.2 million users and 29 million dyadic responses (“like” indications).

Netflix 5 star data. For the sake of reproducibility of our results, we also report results on a data set derived from the Netflix prize data⁵, one of the most famous public data sets for recommendation. The Netflix data set contains 480K users and 18K movies. We derive binary responses by considering only 5-star ratings as “positive” dyads and treating all the others as missing entries.

⁵<http://www.netflixprize.com>




Figure 1: Histograms of the predicted dyadic responses $\{\hat{y}_{ui}\}$: while the predictions by CF (top) are over-optimistically concentrated on positive responses (i.e. predicting “relevant” for all possible dyads), the results obtained by CCF (bottom) demonstrate a more realistic power-law property.

Table 3: Top- k ranking performance on two binary dyadic data sets with simulated contexts.

Model		AP@5	AR@5	nDCG@5
Social				
CF	ℓ_2	0.448	0.230	0.475
CF	Logistic	0.449	0.230	0.476
CCF	Softmax	0.688	0.261	0.704
CCF	Hinge	0.686	0.260	0.702
Netflix-5star				
CF	ℓ_2	0.135	0.022	0.145
CF	Logistic	0.135	0.023	0.146
CCF	Softmax	0.186	0.033	0.189
CCF	Hinge	0.185	0.032	0.188

For both data sets, we randomly split the data into three pieces, one for training, one for testing and the other for validation.

Evaluation metrics. We assess the recommendation performance of each model by comparing the top suggestions of the model to the true actions taken by a user (i.e. “like” or 5-star). We consider three measures commonly used for accessing top- k ranking performance in the IR community:

AP is the *average precision*. AP@ n averages the precision of the top- n ranked list of each query (e.g. user).

AR or *average recall* is the average recall of the top- n rank list of each query.

nDCG or *normalized Discounted Cumulative Gain* is the normalized position-discounted precision score. It gives larger credit to top positions.

For all the three metrics we use $n = 5$ since most social networks and movie recommendation sites recommend a similar number of items for each user visit.

Evaluation protocol. We compare the two CCF models (i.e. Softmax and Hinge) with the two standard CF factorization models (i.e. ℓ_2 and Logistic) described in §2.2. For dyadic data with binary responses, the Logistic CF model amounts to the state-of-the-art [22, 1].

We adopt a fairly strict top- k ranking evaluation. For each user, we assess the top results out of a total preference ordering of the whole item set. In particular, for each user u , we consider all the items as candidates; we compute the three measures based on the comparison between the ground truth (the set of items in the test set that user u actually liked) and the top-5 suggestions predicted by each

model. For statistical consistency, we employ a cross-validation style procedure. We learn the models on training data with parameters tuned on validation data, and then apply the trained models to the test data to assess the performance. All three measures reported are computed on test data only, and they are averaged over five random repeats (i.e. random splits of the data).⁶

To render the data compatible with CCF we simulate a fixed-size pseudo-offer set for each interaction. Specifically, for every positive observation, e.g. $y_{ui} = 1$, we randomly sample a handful set of missing (unobserved) entries $\{y_{ui'}\}_{i'=1:m}$. These sampled dyads are then treated as non-choices, and together with the positive dyad, they are used as the offer set for the current session. In our experiments, we choose $m = 9$ pseudo non-choices; in other words, we assume the offer size $|\mathcal{O}_t| = 10$.

Results and analysis. We report the mean scores in Table 3. Since the dataset are fairly large the standard deviations of all values were below 0.001. Consequently we omitted the latter from the results. As can be seen from the table, CCF dramatically outperforms CF baselines on both data sets. In terms of AP@5, the two CCF models gain about 52.8%–53.6% improvements compared to the two CF models on the **Social** data, and by 37.0%–37.8% on **Netflix-5 star**. Similar comparisons apply to the nDCG@5 measure. And in terms of the AR@5, CCF models outperform CF competitors by up to 13.5% on **Social**, and 30% on **Netflix-5 star** data. All these improvements are statistically highly significant. Note that these results are quite consistent: both CF models perform comparably with each other on both data sets; the performance of the two CCF variants is also comparable; between the two groups, there are noticeable gaps.

One argument we made in this paper for motivating our work is that since the CF models disregard the context information and only learns on positive (action) dyads, they almost inevitably yield overly-optimistic predictions (i.e. predicting positive for all possible dyads). We hypothesize that such estimation *bias* is one of the key reasons for the inability of CF models in learning binary dyadic data. As an empirical validation, in Figure 1, we plot the histograms of the predicted dyadic responses \hat{y}_{ui} (i.e. entries of the diffused matrices) obtained by a CF model (ℓ_2) and a CCF model respectively.⁷ As we can see, the CF model indeed predicts “positive” for most (if not all) dyads; in contrast, the results obtained by the CCF model demonstrate a more realistic power-law distribution [8].⁸

⁶ Note that the contextual information (the offer set \mathcal{O}_t for each interaction t) is missing for both of the two dyadic data sets. We choose the datasets anyway to ensure that the results (at least on the Netflix dataset) can be repeated by other research groups. Results on interaction data are reported in §4.2.

⁷ Similar results obtained with other losses.

⁸ Note that the distribution starts at around 0.5 instead of 0, which is consistent with our intuitions

In reality, each user can only afford to “like” a few items out of a huge amount of alternatives. This power-law property is crucial for information filtering because we are intended to identify a few truly relevant items by *filtering* out many many irrelevant ones. A power-law recommender is desirable in a way analogous to a filter with narrow-bandwidth, which effectively filters the noises (i.e. irrelevant items) and only let the true signal (i.e. relevant items) pass to the endnode (i.e. users).

4.2 User-system interaction data

We now move on to a more realistic evaluation by applying CCF to real user-system interaction data. We evaluate CCF in both an offline test and an online test while comparing its results to both CF baselines.

Data. We collected a large-scale set of user-system interaction traces from a commercial article (News feeds) recommender system. In each interaction, the system offers four personalized articles to the visiting user, and the user chooses one of them by clicking to read that article. The recommendations are dynamically changing over time even during the user’s visit. The system regularly logs every click event of every user visit. It also records the articles being presented to users at a series of discrete time points. To obtain a context set for each user-system interaction, we therefore trace back to the closest recording time point right before the user-click, and we use the articles presented at that time point as the offer set for the current session. We collected such interaction traces from logged records of over one month. We use a random subset containing 3.6 million users, 2500 items and over 110 million interaction traces. Learning an effective recommender on this data set is particularly challenging as the article pool is dynamically refreshing, and each article only has a lifetime of several hours — it only appears once within a particular day, is pulled out from the pool afterward never to appear again.

Evaluation protocol. We consider the following two evaluation settings, one of-line and the other offline.

Offline evaluation Similar to the evaluations presented in §4.1, we evaluate the learned recommender models in terms of the top- k ranking performance on a hold-out test subset. We follow the same configurations in §4.1 and use the three ranking measures, i.e. AP@n, AR@n and nDCG@n as the evaluation

that there is actually *no* truly “irrelevant” item for a user – any item has potential utility for a user; user choose one over another based on the relative preference rather than absolute utility. This is true especially in this era of information explosion, where a user is typically facing so many alternatives that she can only pick the one she likes most while ignoring the others.

Table 4: Offline test (top- k ranking performance) on user-system interaction data.

Model		AP@4	AR@4	nDCG@4
30% Training				
CF	ℓ_2	0.245	0.261	0.255
CF	Logistic	0.246	0.263	0.257
CCF	Softmax	0.262	0.278	0.274
CCF	Hinge	0.261	0.278	0.273
50% Training				
CF	ℓ_2	0.250	0.273	0.268
CF	Logistic	0.252	0.276	0.269
CCF	Softmax	0.266	0.285	0.278
CCF	Hinge	0.265	0.285	0.277
70% Training				
CF	ℓ_2	0.253	0.275	0.271
CF	Logistic	0.253	0.276	0.274
CCF	Softmax	0.267	0.287	0.280
CCF	Hinge	0.267	0.286	0.280

metrics. Note that here we use $n = 4$ instead of 5, because it is the default recommendation size used in the news recommender system.

Online evaluation We further conduct an online test. In particular, for each incoming interaction, we use the trained models to predict which item among the four recommendations will be taken by the user. This prediction is of crucial importance because one of the key objectives for a recommender system is to maximize the traffic and monetary revenue by lifting the click-through rate.

Offline test results. In this setting, we train each model on progressive proportions of 30%, 50% and 70% randomly-sampled training data respectively, and evaluate each trained model in terms of offline top- k ranking performance. The results are reported in Table 4. The two CCF models greatly outperform the two CF baselines in all the three evaluation metrics. Specifically, CCF models gain up to 6.9% improvement over the two CF models in terms of average precision; up to 6.5% in terms of average recall, and up to 7.5% in terms of nDCG. We also conducted a t -test with a standard 0.05 significance level. The hypothesis tests indicate that all the improvements obtained by CCF are significant.

It is worth noting that the improvements obtained by CCF compared to CF baselines are especially evident when the training data are sparser (e.g. using only

30% of training data). This observation empirically validates our argument that the contexts contain substantial useful information for learning recommender models especially when the dyadic action responses are scarce.

Table 5: Online test (predicted click probability) on user-system interaction data.

Model	30%train	50%train	70%train
Random	0.250		
CF ℓ_2	0.337	0.343	0.347
CF Logistic	0.341	0.345	0.347
CCF Softmax	0.376	0.384	0.391
CCF Hinge	0.377	0.385	0.391

The offline results obtained by CCF are quite satisfactory. For example, the average precision is up to 0.276, which means, out of the four recommended items, on average 1.1 are truly “relevant” (i.e. actually being clicked by the user). This performance is quite promising especially considering that most of the articles in the content pool are transient and subject to dynamically updating.

Online test results. We further evaluate the online performance of each compared model by assessing the predicted click rates. Click-rate is essential for an online recommender system because it is closely-related to both the traffic and the revenue of a webshop. In our evaluation, for each of the incoming visits $(u_t, \mathcal{O}_t, i_t^*)$, we use the trained models to predict the user choice, i.e. we ask the question: “among all the offered items $i \in \mathcal{O}_t$, which one will most likely be clicked?” We use the trained model to rank the items in the offer set, and compare the top-ranked item with the item that was actually taken (i.e. i_t^*) by user u_t . We evaluate the results in terms of the prediction accuracy.

The results are given in Table 5. Because the size of each offer set in the current data set is 4, a random predictor yields 0.25. As seen from the table, while all the four models obtain significantly better predictions than the random predictor, the two CCF models further greatly outperform the two CF models. Specifically, we observe 11.3%–12.7% improvements obtained by CCF models compared to the two CF competitors. These results are quite significant especially considering the dynamic property of the system.

Impact of parameters. The performance of the two CCF models is affected by the parameter settings of the latent dimensionality, k , as well as the regularization weights, $\lambda_{\mathcal{I}}$ and $\lambda_{\mathcal{U}}$. In Figure 2⁹, we illustrate how the offline top- k ranking per-

⁹Due to heavy computational consumptions, these results are obtained on a relatively small subset of data.




Figure 2: Offline top- k ranking performance (nDCG@5) as a function of latent dimensionality (top) and regularization weight (bottom).

formance changes as a function of these parameters, where we use the same value for both $\lambda_{\mathcal{I}}$ and $\lambda_{\mathcal{U}}$. Here we only reported the results with nDCG@5 measure because the results show similar shapes when other measures (including the click rate) are used. As can be seen from the Figure, the nDCG curves are typically in the inverted U-shape with the optimal values achieved at the middle. In particular, for both the two CCF models, the dimensionality around 10 and regularization weight around 0.0001 yield the best performance, which is also the default parameter setting we used in obtaining our reported results.

Nonresponded sessions. In Section 3.4 we presented two models for encoding nonresponded interactions, e.g. a user visits the News website but does not click any of the recommended articles. These approaches are promising because compared to the responded sessions, the nonresponded ones are typically much more plentiful and if learned successfully, this wealth of information has a potential to alleviating the critical data-sparse issue in recommendation.

Unfortunately, due to the data-logging mechanism of the News recommender system, we were unable to obtain such nonresponded interactions (this is subject to future work). Instead, for a preliminary test, we conducted evaluation on a small set of *pseudo* nonresponded sessions that are derived from the responded ones. In particular, we hold out a randomly-sampled subset of sessions; for each of these sessions, we hide the item being clicked by the user, and use the remaining items as a nonresponded context set by assuming no click for this set. We augmented this set of derived nonresponded sessions to the training set, and train the model on the combined training data. The results from this preliminary evaluation did not show significant performance improvement. This is likely due to the fact that the surrogate distribution is invalid. A detailed analysis with more realistic data is the subject of future research.

5 Related Work

Although a natural reflection of a user’s preference is the process of interaction with the recommender, to our knowledge, this interaction data has not been exploited for learning recommender models. Instead, research on recommender systems has focused almost exclusively on learning the dyadic data. Particularly collaborative filtering approaches only capture the user-item dyadic data with explicit user actions while the context dyads are typically treated missing values. For example, the rating-oriented models aim to approximating the ratings that users assigned to items [25, 20, 24, 1, 6, 15]; whereas the recently proposed ranking-oriented algorithms [29, 16] attempt to recover the ordinal ranking information derived from the ratings.

By exploiting past records of user-item dyadic responses for future prediction based on either *neighborhood based* [25, 20, 16] or *latent factor based* methods [24, 1, 6, 15, 29], collaborative filtering approaches encode the collaboration effect that similar users get similar preference on similar items. In this paper, by leveraging the user-recommender interaction data, we show that much better recommender performance can be obtained when a local-competition effect underlying the user choice behaviors is also encoded.

The multinomial logit model we present is derived based on the random utility theory [17, 18]. The model is well-established and has been widely used for a long time in, e.g. psychology [17], economics [19, 18] and marketing science [11]. Particularly, [11] applied the model to examine the brand choice of households on grocery data; [10] showed this model is theoretically and empirically superior to the ℓ_2 regression model. More recently, the pioneering work of [9] first applied the model to characterize online choices in recommender system and investigated how recommender systems impact sales diversity. Following these steps, this work further employs the model to learn factorization models for recommendation.

The Hinge formulation of CCF shows close connection to the pairwise preference learning approaches widely used in Web search ranking [12]. Our model, however, differs from these content filtering models [12] in that instead of learning a feature mapping as in [12], our model uses the formulation for learning a multiplicative latent factor model.

6 Summary and future research

We presented a framework for learning recommender by modeling user choice behavior in the user-system interaction process. Instead of modeling only the sparse binary events of user actions as in traditional collaborative filtering, the proposed collaborative-competitive filtering models take into account the contexts in which user decisions are made. We presented two models in this spirit, established efficient learning algorithms and demonstrated the effectiveness of the proposed approaches with extensive experiments on three large-scale real-world recommendation data sets.

There are several promising directions for future research.

Attention budget and position bias. When deriving the CCF model, we admit an assumption that user decides whether to take an offer solely based on the comparison of utilities. This assumption, however, neglects a factor which might be important in practice. In particular, a user might have budgeted attention such that when making choices he only pays attention to a few top-ranked items and totally disregard the others. This *position bias* is evident in both web search ranking and

recommendation. We plan to take this into consideration for building choice models.

Recommender strategy and user behavior. A key feature of the current paper is that we assume the recommender adopts a deterministic strategic policy when making recommendations. In practice, a recommender could also adaptively react to the users' actions as well as its own considerations (e.g. inventory constraints, promotion requirement of certain brands). We would like to extend our analysis here to model the interactive process between users and recommender.

Further empirical validation. Due to data collection constraints, some parts of the proposed models are not strictly evaluated in the current paper. We plan to refine the mechanism for data collection and conduct experiments for further experiments.

References

- [1] D. Agarwal and B.-C. Chen. Regression-based latent factor models. In *15th ACM SIGKDD International conference on Knowledge Discovery and Data Mining*, pages 19–28, 2009.
- [2] E. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing. Mixed membership stochastic blockmodels. In *Advances in Neural Information Processing Systems 20*, pages 33–40, 2008.
- [3] Y. Bakos and E. Brynjolfsson. Bundling and competition on the internet. *Marketing Science*, 19(1):63–82, 2000.
- [4] E. Brynjolfsson, Y. J. Hu, and M. D. Smith. Consumer surplus in the digital economy: Estimating the value of increased product variety at online book-sellers. *Management Science*, 49(11):1580–1596, 2003.
- [5] D. E. Byrne. *The attraction paradigm*. Academic Press, 1971.
- [6] Y. Chen, D. Pavlov, and J. F. Canny. Large-scale behavioral targeting. In *15th ACM SIGKDD international conference on Knowledge Discovery and Data Mining*, pages 209–218, 2009.
- [7] W. Chu and S.-T. Park. Personalized recommendation on dynamic content using predictive bilinear models. In *WWW'09: Proceedings of the 18th international conference on World wide web*, pages 691–700, 2009.

- [8] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM'99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*, pages 251–262, 1999.
- [9] D. M. Fleder and K. Hosanagar. Recommender systems and their impact on sales diversity. In *EC'07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 192–199, 2007.
- [10] D. H. Gensch and W. W. Recker. The multinomial, multiattribute logit choice model. *Journal of Marketing Research* Vol. 16, No. 1, pp. 124–132. Feb. 1979.
- [11] P. M. Guadagni and J. D. Little. A logit model of brand choice calibrated on scanner data. *Marketing Science*, 2(3):203–238, 1983.
- [12] R. Herbrich, T. Graepel, and K. Obermayer. Support vector learning for ordinal regression. In *International Conference on Artificial Neural Networks*, pages 97–102, 1999.
- [13] O. Kallenberg. *Probabilistic symmetries and invariance principles*. Springer, 2005.
- [14] Y. Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *KDD'08: Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 426–434, 2008.
- [15] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [16] N. N. Liu and Q. Yang. Eigenrank: a ranking-oriented approach to collaborative filtering. In *SIGIR'08: Proceedings of the 31st ACM SIGIR conference on Research and development in information retrieval*, pages 83–90, 2008.
- [17] R. D. Luce. *Individual choice behavior*. Wiley, 1959.
- [18] C. F. Manski. Maximum score estimation of the stochastic utility model of choice. *Journal of Econometrics*, (3):205–228, August 1975.
- [19] D. McFadden. Conditional logic analysis of qualitative choice behavior. In *Frontiers of econometrics*, Academic Press, 1974.

- [20] M. R. McLaughlin and J. L. Herlocker. A collaborative filtering algorithm and evaluation metric that accurately model the user experience. In *SIGIR'04: Proceedings of the 27th ACM SIGIR conference on Research and development in information retrieval*, pages 329–336, 2004.
- [21] M. McPherson, L. S. Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.
- [22] K. Miller, T. Griffiths, and M. Jordan. Nonparametric latent feature models for link prediction. In *NIPS'09: Advances in Neural Information Processing Systems 22*, pages 1276–1284, 2009.
- [23] B. Murthi and S. Sarkar. The Role of the Management Sciences in Research on Personalization. *Management Science*, 49(10):1344–1362, October 2003.
- [24] R. Salakhutdinov and A. Mnih. Bayesian probabilistic matrix factorization using markov chain monte carlo. In *ICML'08: Proceedings of the 25th international conference on Machine learning*, pages 880–887, 2008.
- [25] B. Sarwar, G. Karypis, J. Konstan, and J. Reidl. Item-based collaborative filtering recommendation algorithms. In *WWW'01: Proceedings of the 10th international conference on World Wide Web*, pages 285–295, 2001.
- [26] A. Singh and G. Gordon. A unified view of matrix factorization models. In W. Daelemans, B. Goethals, and K. Morik, editors, *ECML'08: European Conference on Machine Learning*, pages 358–373. Springer, 2008.
- [27] J. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proc. Intl. Conf. Machine Learning*, 2005.
- [28] T. F. Tan and S. Netessine. Is tom cruise threatened? using netflix prize data to examine the long tail of electronic commerce. *Working paper 1361, Wharton school, University of Pennsylvania*, 2010.
- [29] M. Weimer, A. Karatzoglou, Q. Le, and A. Smola. Cofi rank - maximum margin matrix factorization for collaborative ranking. In *NIPS'07: Advances in Neural Information Processing Systems 20*, pages 1593–1600, 2007.
- [30] K. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML'09: Proceedings of the 26th International Conference on Machine Learning*, pages 1113–1120, 2009. ACM.

- [31] Z. Zheng, H. Zha, T. Zhang, O. Chapelle, K. Chen, and G. Sun. A General Boosting Method and its Application to Learning Ranking Functions for Web Search. In *NIPS'08: Advances in Neural Information Processing Systems 20*.
- [32] M. Zinkevich, M. Weimer, A. Smola, and L. Li. Parallelized stochastic gradient descent. In *Advances in Neural Information Processing Systems 23*, 2010.